CCC and CRA's Response to the Office of the Assistant Secretary for Research and Technology [Request for Information-Research Ideas To Support Nationwide Automated Vehicle (AV) Deployment]

October 17, 2025

Written by: Gabrielle Allen (University of Wyoming), Haley Griffin (Computing Community Consortium/Computing Research Association), Ming Lin (University of Maryland), Manish Parashar (University of Utah), and Weisong Shi (University of Delaware)

This response is prepared by the Computing Community Consortium (CCC), in collaboration with the Computing Research Association (CRA). CRA is an association of over 270 North American computing research organizations, both academic and industrial, and partners from six professional computing societies. CCC's mission, a CRA subcommittee, is to enable the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges.

Please note any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the authors' affiliations, or of the U.S. National Science Foundation (NSF), which funds the CCC.

Computing Community Consortium (CCC) and Computing Research Association (CRA) appreciate the opportunity to provide input to the U.S. Department of Transportation's (U.S. DOT) Office of the Assistant Secretary for Research and Technology (OST-R) regarding your Request for Information on coordinated national research supporting Automated Driving Systems (ADS) transportation technology deployment. We commend OST-R's proactive approach in exploring how ADS-equipped vehicles, which will change the lives of many Americans, can be deployed in a safe, data-driven manner.

1. Data Standards and Integration

Research is needed to establish comprehensive data frameworks that can standardize information on autonomous vehicles (AV) interactions and their impact on transportation systems. These frameworks are essential for categorizing, reporting, and validating data related to both normal and anomalous events, such as obstructions or unexpected AV behaviors. Central to this is developing an API and vehicle interface that allows seamless, standardized data exchange with infrastructure. It is key to collaborate with industry partners (Waymo, Google, Uber, Lyft, etc.) to access current ADS transportation data at scale. Specifically, we need research into data standards and interoperability models to enable the creation of national adverse-weather datasets that can support algorithm development.

1(a) Statistical Methodologies for Benchmarking

There is a need for new or enhanced statistical methodologies to create reliable benchmarks for AV impacts on transportation system performance. This includes applying new metrics, confidence intervals, as well as considering appropriate human driver baselines (or "negative control events"). It is also important to consider new benchmarking measures in key areas like safety, energy consumption, and overall system (e.g. road network) utility (e.g. flows, capacities, units served).

1(b) and 2(a) Specific Data and Publication Methods: Edge Cases

Edge-case collection and sharing should be a national flagship investment, because long-tail failures—not averages—are what stall trust and nationwide scaling. Two concrete, recent cases show why: a Cruise pedestrian-dragging incident in San Francisco led to a permit suspension and a federal settlement (DMV suspension statement; U.S. DOJ deferred prosecution announcement), and a Waymo robotaxi struck a cyclist after an occlusion near 17th & Mississippi in SF—an archetypal "rare but foreseeable" urban scenario (SF Standard report; ABC7 coverage). In parallel, the National Highway Traffic Safety Administration's (NHTSA) probe of Waymo has documented repeated crashes with stationary or semi-stationary objects like gates and chains, signaling systematic blind spots that different companies keep rediscovering (PE24-016 investigation memo). These show that proprietary incident handling yields slow, uneven learning. By contrast, a federally backed, standardized edge-case program and novel research methodologies (e.g. use of digital twin technology, real-to-sim data analysis on edge-cases) would compress the safety-learning cycle across all Operational Design Domains (ODDs).

National Edge-Case Commons

We recommend creating a National Edge-Case Commons that includes (1) a shared taxonomy and schema (event/near-miss triggers, ODD/context, object types, work-zone/map flags); (2) privacy-preserving uploads of redacted sensor snippets + decision traces into a write-once/verify-always repository with clear chain-of-custody; (3) open APIs and "negative-control" human-driver baselines; and (4) a public tier integrated with DOT's open portal so agencies and researchers can audit rare-event performance uniformly across fleets and geographies—e.g., integration points with the ITS DataHub (portal). This directly answers the

RFI's calls on standardized information exchange, longitudinal tracking, and consistent edge-case mapping—turning one operator's hard-won lesson into everyone's safety gain, faster.

A Nationwide Edge-Case & Rare-Event Registry would make roads safer by finding what today's systems miss: unusual, high-risk situations at intersections. Using OpenIntersection (sensor-equipped, software-defined intersections) and an Intersection Programming Interface (IPI) (a common, open API), local infrastructure continuously flags near-misses, occlusions, wrong-way movements, and other anomalies across urban, suburban, and rural sites. Each event is standardized, privacy-protected, and cryptographically verifiable, then shared through open datasets and dashboards so vehicle developers, cities, and researchers can act on a single, evidence-grade picture of risk. The registry enables apples-to-apples evaluation of vehicle behavior (car-following, lane changes, and pedestrian/cyclist interactions) under weather and work-zone stress, driving targeted fixes and policy. The deliverables would be a national taxonomy of rare events, a secure ingestion and labeling pipeline, interoperable APIs, and repeatable test protocols—so safety gains can be measured and scaled.

A shared national registry of edge cases is particularly crucial given that not all companies or regions can conduct comprehensive testing across all challenging conditions (particularly for poor weather conditions like snow, fog, sandstorms, thunderstorms, and heavy rain). Identifying characteristics of scenarios leading to accidents—and collecting that kind of "missing data"—is particularly challenging.

Resilient Edge AI and Federated/Distributed Learning for Low-Connectivity Corridors

Many rural corridors lack reliable broadband or cellular connectivity. Research is needed on distributed AI architectures that enable real-time safety decisions without continuous cloud access. Example important research questions include: What architectures allow low-latency, safety-critical inference at the edge, suitable for deployment in vehicles and roadside stations. How can federated or split learning approaches be adapted for fleets (e.g., snowplows, freight vehicles) to share model improvements without centralized data aggregation? How can solutions be optimized for bandwidth constraints, privacy, and real-time performance in extreme weather?

In addition, micromobility systems are increasingly present across all urban environments and campuses around the US and across the globe. They will also need to be taken into consideration in designing civil infrastructures, public and safety policies, highway and road network design.

1 (d) Infrastructure-Related Data Needs and Exchange

Data Infrastructure and Standardization

Edge cases and challenging operating conditions must be explicitly prioritized in data infrastructure development. New methods are needed to integrate real-time sensor data with satellite imagery and GPS to maintain localization when landmarks and road markings are degraded, obscured or absent. As important are standardized frameworks for collecting, labeling, and sharing adverse weather and other challenging condition datasets at scale. Such frameworks would enable reproducible research and accelerate technology transfer.

National Software-Over-The-Air (SOTA) Update Framework

Near-term infrastructure-related data needs also include supporting large-scale deployment and software updates, particularly for machine learning (ML) models. There is a demand for better infrastructure like a National Software-Over-The-Air (SOTA) Update Framework. SOTA updates are remote, wireless updates delivered via cellular or Wi-Fi networks; they allow manufacturers to push new features, map data, AI-driven improvements, and safety fixes to vehicles without a dealership visit. To ensure these updates are safe and interoperable, the program would define common data schemas and APIs for update telemetry—covering pre-update risk assessments (e.g., SBOM-linked vulnerabilities), version metadata, ODD-slice targeting, and controlled canary/rollback strategies. Building on the RFI's emphasis on data integrity and chain-of-custody, the framework would employ cryptographically verifiable logs and write-once/read-many storage to preserve evidence across in-vehicle, edge, and central repositories. It also aligns with research noting that AVs must regularly receive OTA updates to address bugs and security vulnerabilities, yet those updates introduce significant cyber-security risks; thus, rigorous verification, authentication and end-to-end encryption are essential.

A companion research thrust would establish multi-site SOTA testbeds and an evaluation framework to assess the performance, safety and resilience of update pipelines across varied environments. These testbeds would integrate roadside compute (e.g., 5G/MEC), standardized digital work-zone and map feeds, and tamper-evident update logs to meet the Evidence Act's requirements. They would evaluate how updates impact AI module performance (e.g., object detection or path planning) and vehicle behavior consistency—car-following, lane changes, and vulnerable road-user detection—under diverse weather, work-zone and connectivity conditions. The evaluation framework should include metrics, confidence intervals, and negative-control events to provide statistically sound evidence of system-level effects. By coordinating federal testbeds and publishing results through DOT-supported research hubs and open-data portals, this program will create a replicable national benchmark for SOTA updates. Deliverables will include: (1) a reference SOTA data and API specification; (2) standardized evaluation metrics for pre- and post-update performance; and (3) guidelines for public reporting and data sharing. Together, these actions ensure that SOTA updates continually enhance AV capabilities—improving AI functions, safety features, and map accuracy—while upholding safety, security, equity and public trust across the U.S. transportation system.

3. Supervision Dynamics and Human Interaction

(a, c) Research required for long-term transportation planning with AVs centers on the mixed autonomy environment where human and robot drivers coexist. Key research areas include understanding and characterizing the dynamics between pedestrians, human drivers, passengers and autonomous vehicles —an HCI and computing problem—to ensure safety. There is a particular need for more research about AV and human interactions at intersections, which are high-risk areas, and during poor weather conditions (snow, hail, rain, etc.). This research must investigate how AVs interact with human behaviors and how to resolve issues like non-verbal communication and right-of-way protocols, given that current AV actions are considered less predictable than human drivers. There is also a need to study the impact of AVs on overall safety in this mixed pattern and to consider interactions with micromobility systems (e.g. scooters, bicycles, unicycles, etc.) and animals (e.g. dogs, cats, deer, etc.). This research must balance privacy and openness by determining how to make relevant data available from AV operators while protecting proprietary information.

(b) Research is needed to leverage the Evidence Act's requirements for explainability and data provenance; a robust, core infrastructure for logging and storing data. Every time an AV makes a decision, data should be collected so there is a data trail in case of legal probes. AV decision-making needs to be explainable so that in case of an accident, a reconstruction can occur to determine the cause and assess if it is a systematic mistake.

4. Evidence Based Evaluation

(a) There is a need for more research to support safe, transparent, and equitable nationwide evidence-based evaluation of AV operational impacts. This necessitates a wide range of geographical data and consideration of different weather conditions. The establishment of a trusted national repository through a standardized AV Data Storage Standard and Access API (see 5a below) would provide the necessary backbone for consistent, evidence-based evaluation across all jurisdictions.

5. Transparency and Building Public Understanding

(a) Robust data storage and management infrastructure is essential to the safe and scalable deployment of automated vehicles. A standardized, secure, and interoperable storage framework enables the preservation of multimodal sensor and algorithm outputs, providing the foundation for longitudinal research, benchmarking, and evidence-based evaluation of AV performance. Storage systems ensure compliance with data integrity and chain-of-custody requirements, supporting legal admissibility and public accountability while also facilitating reproducible research across diverse operating environments. Without reliable storage, other priorities such as nationwide mapping, fleet analytics, and interoperability cannot be realized at scale.

Federal investment in AV data storage will address critical gaps by providing a unified “backbone” for normalizing, archiving, and sharing AV data across jurisdictions and operators. This infrastructure would reduce fragmentation, enable consistent standards and APIs, and ensure the capture of rare events and pre-crash data essential for safety evaluation. Prioritizing data storage research and deployment amplifies the effectiveness of all other AV initiatives, creating a trusted national repository that strengthens transparency, fosters public confidence, and accelerates the safe integration of automated vehicles into the U.S. transportation system.

(b) There is a need for statistical methods to capture emerging anomalous behavior and rare-event factors associated with AV impacts. A system must be able to incorporate lessons learned along the way, which requires an incentive structure for companies to share data and lessons learned about what does not work so they can be forwarded into research.

(c) There needs to be a national definition of a road in order to properly map the country, and different types of roads as well (public/private, seasonal usage, degree of pavement). This is especially important for rural areas where current navigational software consistently provides incorrect information, which could be incredibly dangerous for an AV.

6. Evaluation of Consistent and Robust Vehicle Behavior Interactions

Challenges Posed by Non-Traditional Vehicles

When developing standardized methods for evaluating AV behavior, it is essential to broaden the scope beyond conventional road vehicles. Smaller, non-traditional AVs, such as delivery robots, powered wheelchairs, and electric scooters and skateboards, introduce unique operational challenges and expose policy gaps. These non-traditional vehicles often operate in shared spaces like sidewalks and bike lanes where there is a lack of AV rules and standards to guide right of way and interaction protocols with other AVs, pedestrians, animals etc. Robust modeling and simulation is required to predict human behavior and limitations especially in the context of pedestrians on both AV and non AV-wheels (bikes, scooters, rollerblades, skateboards, etc.). Expanding standardized operational frameworks that predict human behavior in these mixed environments is critical for developing evaluation methods that reflect real-world conditions and ensuring safety in increasing diverse transportation ecosystems.

Physics-Informed, Multi-Sensor AI and Intelligent Roadside Infrastructure for Adverse Conditions and Rural Areas

Most AV perception and control research has been developed and benchmarked in urban, temperate settings, leaving a major gap in sensor fusion algorithms and predictive modeling for adverse weather conditions such as snow, ice, and high winds. These challenges could be addressed by integrating thermal, radar and LiDAR data to detect hazards invisible to traditional sensors such as black ice and developing physics-informed AI models that can rapidly predict changes in road condition to mitigate risks like loss of traction and vehicle blow-over. Progress in this area will require new benchmark datasets and simulation frameworks to represent adverse conditions for AV perception and control research.

Multi-modal roadside infrastructure could play an important role in enhancing AV performance under adverse conditions. Current localization and perception methods degrade severely under snow cover, GNSS failure, and white-out conditions. Foundational research is needed on how to design and deploy intelligent roadside sensing systems that provide robust, high-fidelity environmental information to provide “beyond-the-horizon” awareness using the combined modalities of LiDAR, radar, thermal/optical cameras, and weather sensors in real time? New methods would also be needed to integrate data from these systems into digital twins to model and predict evolving hazards under extreme weather conditions.